

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant(s):	§	Customer No.:	29855
Venkat Rangan	§		
Ed McClanahan	§		
Michael Schmitz	§	Art Unit:	2182
Serial No.: 10/695,408	§	Confirmation No.	5639
Filed: October 28, 2003	§	Examiner:	Scott C. Sun
For: APPARATUS AND METHOD FOR DATA MIGRATION IN A STORAGE PROCESSING DEVICE	§	Docket No.	112-0122US

Box Appeal Brief  
Commissioner for Patents  
P. O. Box 1450  
Alexandria, VA 22313-1450

**Mail Stop: Appeal Briefs – Patents**

## **APPEAL BRIEF**

**TABLE OF CONTENTS**

I.	REAL PARTY IN INTEREST.....	3
II.	RELATED APPEALS AND INTERFERENCES.....	4
III.	STATUS OF CLAIMS .....	5
IV.	STATUS OF AMENDMENTS .....	6
V.	SUMMARY OF CLAIMED SUBJECT MATTER .....	7
VI.	GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL .....	10
VII.	ARGUMENT.....	11
A.	§ 103 Rejections.....	11
B.	Claim 19.....	11
C.	Claim 21 .....	13
D.	Remaining Claims.....	15
E.	Conclusion .....	15
VIII.	CLAIMS APPENDIX.....	17
IX.	EVIDENCE APPENDIX.....	24
X.	RELATED PROCEEDINGS APPENDIX.....	25

**I. REAL PARTY IN INTEREST**

Brocade Communications Systems, Inc. is the real party in interest

**II. RELATED APPEALS AND INTERFERENCES**

None

**III. STATUS OF CLAIMS**

Claims 1-36 are rejected. The appealed claims are 1-36.

**IV. STATUS OF AMENDMENTS**

None filed

**V. SUMMARY OF CLAIMED SUBJECT MATTER**

This section provides a concise explanation of the subject matter defined in each of the independent claims involved in the appeal, referring to the specification by paragraph and line number and to the drawings by reference characters as required by 37 CFR § 41.37(c)(l)(v). Each element of the claims is identified with a corresponding reference to the specification and drawings where applicable Note that the citation to passages in the specification and drawings for each claim element does not imply that the limitations from the specification and drawings should be read into the corresponding claim element.

Independent claim 1 provides for a storage processing device (SAP 100, Figs. 1A, 1B; ¶ 40) for use in a switched fabric (network 114, Fig. 1A; fabric 116, Fig. 1B; ¶ 41), the switched fabric including at least one switch (¶ 41) and the storage processing device, with the storage processing device coupled to the at least one switch (¶ 44), with first and second storage units (iSCSI target 112, FC target 106, Figs. 1A, 1B; ¶¶ 40-44) connected to and coupled through the switched fabric and coupled to the at least one switch and the storage processing device, where the first and second storage units may be directly connected to the storage processing device or may be coupled through the at least one switch (Figs. 1A, 1B; ¶ 44), the storage processing device to migrate data from the first storage unit to the second storage unit (Figs. 4, 6, 19, 20A, 20B; ¶¶ 215-220), whether the first and second storage units are directly connected to the storage processing device or are coupled through the at least one switch the storage processing device comprising:

an input/output module (200, Figs. 2, 4; ¶¶ 46, 52, 53) for coupling to the first and second storage units including processors (400, Figs. 4, 5; ¶¶ 52, 55-63) to receive, operate on and transmit network traffic; and

a control module (202, Figs. 2, 6; ¶¶ 46, 47, 64) coupled to said input/output module, said input/output module (420, Fig. 4) and said control module (420, Fig. 6) being configured to interactively perform data migration (¶ 215) from the first storage unit to the second storage unit.

Independent claim 10 provides for a switched fabric (network 114, Fig. 1A; fabric 116, Fig. 1B; ¶ 41) for connection to and coupling of at least one host (FC host 102, iSCSI host 108, Figs. 1A, 1B; ¶¶ 40-44) and at least two storage units (iSCSI target 112, FC target 106, Figs. 1A, 1B; ¶¶ 40-44), the fabric comprising:

at least one switch (¶ 41) for coupling to the at least one host and the at least two storage units; and

a storage processing device (SAP 100, Figs. 1A, 1B; ¶ 40) coupled to the at least one switch (¶ 44) and for coupling to the at least one host and first and second storage units of the at least two storage units, where the first and second storage units may be directly connected to the storage processing device or may be coupled through the at least one switch (Figs. 1A, 1B; ¶ 44), the storage processing device to migrate data between first and second storage units of the at least two storage units (Figs. 4, 6, 19, 20A, 20B; ¶¶ 215-220) whether the first and second storage units are directly connected to the storage processing device or are coupled through the at least one switch, the storage processing device including:

an input/output module (200, Figs. 2, 4; ¶¶ 46, 52, 53) including processors (400, Figs. 4, 5; ¶¶ 52, 55-63) to receive, operate on and transmit network traffic; and

a control module (202, Figs. 2, 6; ¶¶ 46, 47, 64) coupled to said input/output module, said input/output module (420, Fig. 4) and said control module (420, Fig. 6) being configured to interactively perform data migration between the first and second storage units.

Independent claim 19 provides for a network comprising:

at least one host (FC host 102, iSCSI host 108, Figs. 1A, 1B; ¶¶ 40-44) adapted to be connected to a switched fabric;

at least two storage units (iSCSI target 112, FC target 106, Figs. 1A, 1B; ¶¶ 40-44), each adapted to be connected to a switched fabric; and

a switched fabric (network 114, Fig. 1A; fabric 116, Fig. 1B; ¶ 41) connected to and coupling the at least one host and the at least two storage units, the switched fabric comprising:

at least one switch (¶ 41) for coupling to the at least one host and the at least two storage units; and

a storage processing device (SAP 100, Figs. 1A, 1B; ¶ 40) coupled to the at least one switch (¶ 44) and for coupling to the at least one host and first and second storage units of the at least two storage units, where the first and second storage units may be directly connected to the storage processing device or may be coupled through the at least one switch (Figs. 1A, 1B; ¶ 44), the storage processing device to migrate data between the first and second storage units (Figs. 4, 6, 19, 20A, 20B; ¶¶ 215-220) whether the first and second storage units are directly connected to the storage processing device or are coupled through the at least one switch, the storage processing device including:

an input/output module (200, Figs. 2, 4; ¶¶ 46, 52, 53) including processors (400, Figs. 4, 5; ¶¶ 52, 55-63) to receive, operate on and transmit network traffic; and  
a control module (202, Figs. 2, 6; ¶¶ 46, 47, 64) coupled to said input/output module, said input/output module (420, Fig. 4) and said control module (420, Fig. 6) being configured to interactively perform data migration between the first and second storage units.

Independent claim 28 provides for a method for supporting data migration between first and second storage units (iSCSI target 112, FC target 106, Figs. 1A, 1B; ¶¶ 40-44) connected to and coupled through a switched fabric (network 114, Fig. 1A; fabric 116, Fig. 1B; ¶ 41), the switched fabric including at least one switch (¶ 41) and a storage processing device (SAP 100, Figs. 1A, 1B; ¶ 40), the storage processing device coupled to the at least one switch (¶ 44), where the first and second storage units may be directly connected to the storage processing device or may be coupled through the at least one switch (Figs. 1A, 1B; ¶ 44), the method comprising in the storage processing device:

providing an input/output module (200, Figs. 2, 4; ¶¶ 46, 52, 53) including processors (400, Figs. 4, 5; ¶¶ 52, 55-63) receiving, operating on and transmitting network traffic; and

providing a control module (202, Figs. 2, 6; ¶¶ 46, 47, 64) coupled to said input/output module, said input/output module (420, Fig. 4) and said control module (420, Fig. 6) being configured to interactively perform data migration between the first and second storage units whether the first and second storage units are directly connected to the storage processing device or are coupled through the at least one switch.

**VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL**

Claims 1-36 stand rejected under 35 U.S.C. § 103 as being obvious over Edsall, U.S. Patent Appn. 2003/0172149, in view of Testardi, U.S. Patent Appln. 2003/0140210.

**VII. ARGUMENT**

The claims do not stand or fall together. Instead, appellants present separate arguments for various independent and dependent claims. After a concise discussion of cited art, each of these arguments is separately argued below and presented with separate headings and sub-headings as required by 37 CFR § 41.37(c)(l)(vii).

**A. § 103 Rejections**

Applicants respectfully submit that the present claims are allowable.

**B. Claim 19**

Claim 19 is being treated as an exemplary claim in these remarks to better track the Final Office Action. The Final Office Action rejected claim 19 over Edsall and Testardi. Applicants respectfully traverse the rejection.

The Final Office Action provided comments addressing Applicants remarks in a response of July 19, 2006 and revised the rejection to correspond to those comments. An Advisory Action provided comments to Applicants remarks in a response of December 1, 2006. Applicants submit that the claim is patentable even given the various comments and new rejection provided in the Final Office Action.

The rejection apparently somehow combines the basic data transfer capabilities of Edsall (connectors 302 and 204, MAC 304 and 322 and switches 306 and 320) and the fast path of Testardi, though how the two are combined is not stated. In fact, the Final Office Action defines the I/O modules to be the fast path of Testardi without any reference to Edsall. The Edsall pieces do not include the required processors in the I/O module, with the Testardi fast path apparently filling this requirement.

The Advisory Action addresses this argument by stating that the various logic elements (302, 304, 306, 320, 322, 324) are all the required processors. Applicants first note that elements 302 and 304 are connectors. Edsall ¶ 0060. Clearly connectors cannot be processors. Elements 306 and 320 are switches. Edsall ¶ 0053. Clearly switches as used in Edsall are not processors.

Elements 304 and 322 are Media Access Control (MAC) blocks. Edsall ¶¶ 0053, 0059. Clearly MAC blocks are not processors.

The Advisory Action goes on to state that Edsall teaches that a “multi-port switch will have virtualization logic separately implemented on one or more of its ports . . . rather than having centralized processing for all ports of a switch.” While this may be a correct statement, none of the elements 302, 304, 306, 320, 322 or 324 do any such processing. Therefore Applicants continue to submit that Edsall does not teach the required processors, with only the Testardi fast path apparently filling this requirement.

Applicants note that the claim specifically requires that the storage processing device may have the storage units directly connected or coupled through an external switch and the data migration is performed in either case. Applicants submit that the combination of Edsall and Testardi do not teach or suggest this requirement. The switches 306 and 320 of Edsall have been corresponded to being part of the I/O module, not the external switch required by the claims. Testardi is apparently cited for the processing in the I/O module as well as the control module with the required interactivity. However Testardi does not teach or suggest that the storage devices can be connected through a switch and have the processing of Testardi still be operational. Testardi indicates that the storage units must be directly connected to the storage processing device. Figure 4A of Testardi shows direct connection but at a general level. Figure 4B of Testardi is more specific and shows the hosts and storage units directly connected. Figure 4C adds external switches to Figure 4B but only hosts are attached to the switches, not storage units. There is no teaching or suggestion in Testardi that the storage units can be connected to the storage processing device through a switch and have the data migration operations still be successfully performed. While Edsall may show storage units connected to switches, this does not overcome the limitations of Testardi in this regard as the elements being used in Edsall do not contain any of the processing capabilities required in the claim and those processing capabilities are limited to directly connected storage units. Asserting that the storage units could be connected through a switch goes against the express teachings of Testardi and so is an improper rejection.

The Advisory Action first addresses the “may” language in the claim. The Advisory Action argues that the “may” creates alternative language so that if either condition is met, the

rejection is proper. This misconstrues the language of the claim. Applicants reproduce the relevant language here for easier reference.

a storage processing device coupled to the at least one switch and for coupling to the at least one host and first and second storage units of the at least two storage units, where the first and second storage units may be directly connected to the storage processing device or may be coupled through the at least one switch, the storage processing device to migrate data between the first and second storage units whether the first and second storage units are directly connected to the storage processing device or are coupled through the at least one switch

From a reading of this language it is very clear that the “may” is not an alternative as to the operation of storage processing device. It is specifically required to operate in both conditions, not one. Thus the alternative argument made in the Advisory Action is misplaced.

The Advisory Action, helpfully, does go on to address the both conditions situation. The Advisory Action cites Testardi ¶ 56 which states “some or all of the connections by which the hosts, data management system 16 and data storage system 12 may be connected to the connection medium 18 may pass through other communication devices, such as a Fibre Channel switch . . .” While this is an accurate statement, it is inapposite to the situation at hand because the claims relate only to elements entirely inside the data storage system 12 of Testardi. The connections between the hosts and the data storage system are not at issue, except to note the citation only reinforces Applicants position on Figure 4C teaching that hosts can be attached to switches but not storage units.

Applicants thus submit that the combination of Edsall and Testardi, even as combined in the rejection of the Final Office Action and in view of the Advisory Action remarks, do not teach or suggest all of the required claim elements.

Because this required element of claim 19, and similar claims 1, 10 and 28, is missing, Applicants submit that all of the present claims are allowable.

### C. Claim 21

As an exemplary claim, Applicants will address the rejection of claim 21. The Final Office Action stated that Testardi discloses the fast path delaying data write operations if the barrier entry relates to data write operation by faulting the operations to the control path. The Final Office Action then notes that this is interpreted as the fast path delaying the write

operations on the barrier range. Applicants respectfully traverse this interpretation. The claim requires the processors, which are located in the input/output module, to delay the write operations. The cited operation of Testardi has the fast path, equated to the processors of the claim, not handling the operation at all but transferring the operation to the control path, which then faults the operation to force a retry. But the claim requires the processors, i.e. the fast path, delaying the operation. In Testardi the fast path does nothing, not delay, not handle, nothing. It simply passes the operation on to the control path. This does not meet the positive requirement in the claim that the processors perform the delaying operation. Applicants understand and acknowledge the statement that the limitations of the specification are not to be read into the claims but submit that the requirement that the processors delay the write operation is a positive element in the claims and that the fast path of Testardi passing off the operation does not meet the positive requirement of the processors delaying the operation. There is no teaching or suggestion to have the fast path perform a delay rather than pass off the operation.

The Advisory Action attempts to address this argument but apparently misinterprets it. The Advisory Action argues that the Testardi fast path is an implementation of hardware and/or software. Applicants do not dispute this, and, as noted above, have treated the Testardi fast path as being equated to the processors. Applicants instead focus on a required action of the processors, namely delaying write operations if a barrier entry relates to the write operation. Applicants prior statement “the fast path does nothing, not delay, not handle, nothing” was in the context of write operations and barrier entries, not in the abstract as apparently interpreted by the Advisory Action. As Applicants have stated, the claim specifically requires the processors to delay write operations if a barrier entry relates to the write operation. The processors, i.e., fast path for this argument, of Testardi pass off the operation to the control path, which then faults to force a retry. Therefore it is clear that the fast path does not delay write operations if a barrier entry relates to the write operations. The fast path simply does not handle write operations if a barrier entry relates to the write operations. Rephrased again, the fast path takes no active part, i.e., does nothing, for write operations if a barrier entry relates to the write operations. Testardi teaches that such operations must be handled outside of the fast path or processors. This is directly opposed to the claim requirement that the processors delay such write operations.

Applicants submit that claim 21 and similar claims 3, 8, 12, 17, 26, 30 and 35 are allowable.

**D. Remaining Claims**

As all of the remaining claims were rejected on the same grounds as claims 13-15 and like arguments apply, Applicants submit that those claims are allowable.

**E. Conclusion**

For the reasons stated above, Applicants respectfully submit that the rejections should be reversed. Applicants believe that they have complied with each requirement for an appeal brief. If any member of the Board of Appeals has any questions or otherwise feels it would be advantageous, he or she is encouraged to telephone the undersigned at (832) 446-2405.

In the course of the foregoing discussions, Applicants may have at times referred to claim limitations in shorthand fashion, or may have focused on a particular claim element. This discussion should not be interpreted to mean that the other limitations can be ignored or dismissed. The claims must be viewed as a whole, and each limitation of the claims must be considered when determining the patentability of the claims. Moreover, it should be understood that there may be other distinctions between the claims and the prior art which have yet to be raised, but which may be raised in the future.

Application No. 10/695,408  
Appeal Brief

It is believed that a one month extension is necessary and that appropriate fees are due. The Office is authorized to charge such time extension fees per 37 CFR 1.17(a) and any other fees or credit any overpayment to Deposit Account 501922, referencing attorney docket number 112-0122US. If any fees are inadvertently omitted or if any additional fees are required or have been overpaid, please appropriately charge or credit those fees to Wong, Cabello, Lutsch, Rutherford & Brueculeri, LLP Deposit Account Number 501922, referencing attorney docket number 112-0122US.

Respectfully submitted,

/Keith Lutsch/

May 14, 2007

Filed Electronically

Keith Lutsch, Reg. No. 31,851  
Wong, Cabello, Lutsch,  
Rutherford & Brueculeri, L.L.P.  
20333 SH 249, Suite 600  
Houston, TX 77070  
832-446-2405

### VIII. CLAIMS APPENDIX

1. (Previously Presented) A storage processing device for use in a switched fabric, the switched fabric including at least one switch and the storage processing device, with the storage processing device coupled to the at least one switch, with first and second storage units connected to and coupled through the switched fabric and coupled to the at least one switch and the storage processing device, where the first and second storage units may be directly connected to the storage processing device or may be coupled through the at least one switch, the storage processing device to migrate data from the first storage unit to the second storage unit, whether the first and second storage units are directly connected to the storage processing device or are coupled through the at least one switch the storage processing device comprising:

an input/output module for coupling to the first and second storage units including processors to receive, operate on and transmit network traffic; and

a control module coupled to said input/output module, said input/output module and said control module being configured to interactively perform data migration from the first storage unit to the second storage unit.

2. (Previously Presented) The storage processing device of claim 1, wherein said processors include table information related to data migration and wherein said control module is coupled to said table information to maintain said table information for data migration.

3. (Previously Presented) The storage processing device of claim 2, wherein said table information includes a barrier entry and said processors delay data write operations if said barrier entry relates to said data write operation.

4. (Previously Presented) The storage processing device of claim 2, wherein said table information includes an entry related to the extents in the data migration, said entry defining an extent operation type.

5. (Previously Presented) The storage processing device of claim 4, wherein said table information further includes a legend entry for each extent operation type defining migration operations for the extent.

6. (Original) The storage processing device of claim 5, wherein said table information further includes entries referenced by said legend entry defining physical extent location.

7. (Original) The storage processing device of claim 6, wherein legend entries include entries indicating data not migrated, data migrated and a barrier entry for data being migrated.

8. (Previously Presented) The storage processing device of claim 7, wherein said processors delay data write operations if said barrier entry relates to said data write operation.

9. (Original) The storage processing device of claim 8, wherein said control module provides commands to copy data and places said barrier entry for said data being copied.

10. (Previously Presented) A switched fabric for connection to and coupling of at least one host and at least two storage units, the fabric comprising:

at least one switch for coupling to the at least one host and the at least two storage units; and

a storage processing device coupled to the at least one switch and for coupling to the at least one host and first and second storage units of the at least two storage units, where the first and second storage units may be directly connected to the storage processing device or may be coupled through the at least one switch, the storage processing device to migrate data between first and second storage units of the at least two storage units whether the first and second storage units are directly connected to the storage processing device or are coupled through the at least one switch, the storage processing device including:

an input/output module including processors to receive, operate on and transmit network traffic; and

a control module coupled to said input/output module, said input/output module and said control module being configured to interactively perform data migration between the first and second storage units.

11. (Previously Presented) The fabric of claim 10, wherein said processors include table information related to data migration and wherein said control module is coupled to said table information to maintain said table information for data migration.

12. (Previously Presented) The fabric of claim 11, wherein said table information includes a barrier entry and said processors delay data write operations if said barrier entry relates to said data write operation.

13. (Previously Presented) The fabric of claim 11, wherein said table information includes an entry related to the extents in the data migration, said entry defining an extent operation type.

14. (Previously Presented) The fabric of claim 13, wherein said table information further includes a legend entry for each extent operation type defining migration operations for the extent.

15. (Original) The fabric of claim 14, wherein said table information further includes entries referenced by said legend entry defining physical extent location.

16. (Original) The fabric of claim 15, wherein legend entries include entries indicating data not migrated, data migrated and a barrier entry for data being migrated.

17. (Previously Presented) The fabric of claim 16, wherein said processors delay data write operations if said barrier entry relates to said data write operation.

18. (Original) The fabric of claim 17, wherein said control module provides commands to copy data and places said barrier entry for said data being copied.

19. (Previously Presented) A network comprising:  
at least one host adapted to be connected to a switched fabric;  
at least two storage units, each adapted to be connected to a switched fabric; and  
a switched fabric connected to and coupling the at least one host and the at least two storage units, the switched fabric comprising:  
at least one switch for coupling to the at least one host and the at least two storage units; and

a storage processing device coupled to the at least one switch and for coupling to the at least one host and first and second storage units of the at least two storage units, where the first and second storage units may be directly connected to the storage processing device or may be coupled through the at least one switch, the storage processing device to migrate data between the first and second storage units whether the first and second storage units are directly connected to the storage processing device or are coupled through the at least one switch, the storage processing device including:

an input/output module including processors to receive, operate on and transmit network traffic; and

a control module coupled to said input/output module, said input/output module and said control module being configured to interactively perform data migration between the first and second storage units.

20. (Previously Presented) The network of claim 19, wherein said processors include table information related to data migration and wherein said control module is coupled to said table information to maintain said table information for data migration.

21. (Previously Presented) The network of claim 20, wherein said table information includes a barrier entry and said processors delay data write operations if said barrier entry relates to said data write operation.

22. (Previously Presented) The network of claim 20, wherein said table information includes an entry related to the extents in the data migration, said entry defining an extent operation type.

23. (Previously Presented) The network of claim 22, wherein said table information further includes a legend entry for each extent operation type defining migration operations for the extent.

24. (Original) The network of claim 23, wherein said table information further includes entries referenced by said legend entry defining physical extent location.

25. (Original) The network of claim 24, wherein legend entries include entries indicating data not migrated, data migrated and a barrier entry for data being migrated.

26. (Previously Presented) The network of claim 25, wherein said processors delay data write operations if said barrier entry relates to said data write operation.

27. (Original) The network of claim 26, wherein said control module provides commands to copy data and places said barrier entry for said data being copied.

28. (Previously Presented) A method for supporting data migration between first and second storage units connected to and coupled through a switched fabric, the switched fabric including at least one switch and a storage processing device, the storage processing device coupled to the at least one switch, where the first and second storage units may be directly connected to the storage processing device or may be coupled through the at least one switch, the method comprising in the storage processing device:

providing an input/output module including processors receiving, operating on and transmitting network traffic; and

providing a control module coupled to said input/output module, said input/output module and said control module being configured to interactively perform data migration between the first and second storage units whether the first and second storage units are directly connected to the storage processing device or are coupled through the at least one switch.

29. (Previously Presented) The method of claim 28, wherein said processors include table information related to data migration and wherein said control module is coupled to said table information to maintain said table information for data migration.

30. (Previously Presented) The method of claim 29, wherein said table information includes a barrier entry and said processors delay data write operations if said barrier entry relates to said data write operation.

31. (Previously Presented) The method of claim 29, wherein said table information includes an entry related to the extents in the data migration, said entry defining an extent operation type.

32. (Previously Presented) The method of claim 31, wherein said table information further includes a legend entry for each extent operation type defining migration operations for the extent.

33. (Original) The method of claim 32, wherein said table information further includes entries referenced by said legend entry defining physical extent location.

34. (Original) The method of claim 33, wherein legend entries include entries indicating data not migrated, data migrated and a barrier entry for data being migrated.

35. (Previously Presented) The method of claim 34, wherein said processors delay data write operations if said barrier entry relates to said data write operation.

36. (Original) The method of claim 35, wherein said control module provides commands to copy data and places said barrier entry for said data being copied.

**IX. EVIDENCE APPENDIX**

None.

**X. RELATED PROCEEDINGS APPENDIX**

None.